



US 20030058878A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2003/0058878 A1
Minnick et al. (43) Pub. Date: Mar. 27, 2003(54) METHOD AND APPARATUS FOR
MINIMIZING SPINLOCKS AND RETAINING
PACKET ORDER IN SYSTEMS UTILIZING
MULTIPLE TRANSMIT QUEUES

Publication Classification

(51) Int. Cl. 7 H04L 12/28; H04L 12/56
(52) U.S. Cl. 370/412; 370/428; 370/429(76) Inventors: Linden Minnick, Portland, OR (US);
Patrick L. Connor, Portland, OR (US)

(57) ABSTRACT

Correspondence Address:

Lance A. Terries

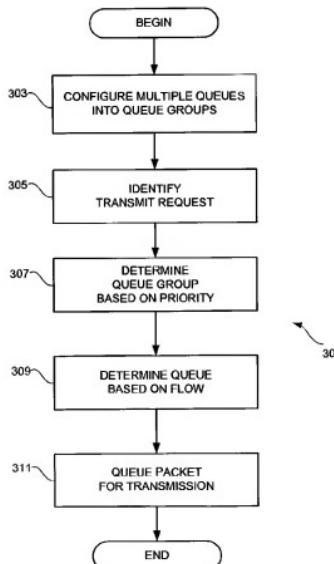
BLAKELEY, SOKOLOFF, TAYLOR & ZAFMAN
LLP
Seventh Floor
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026 (US)

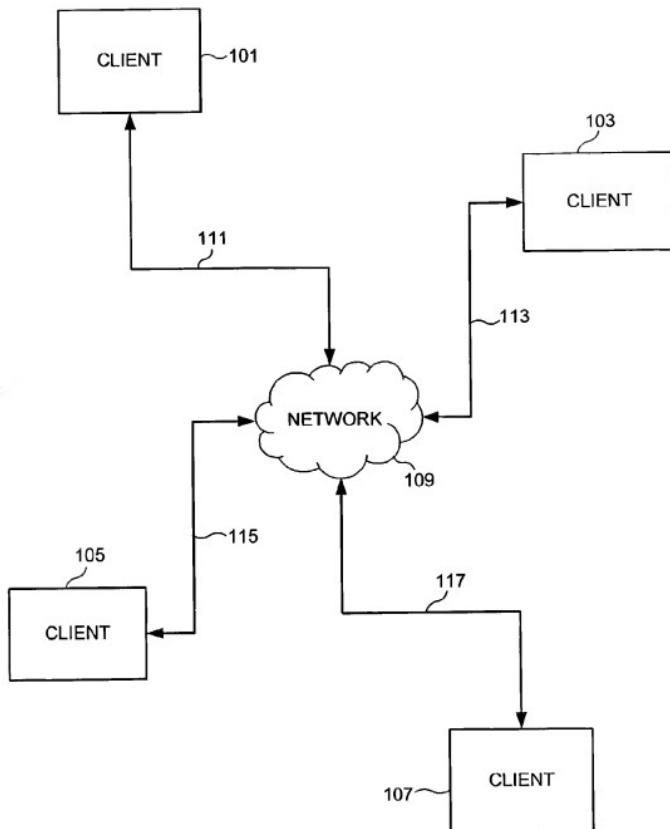
A method, apparatus, and article of manufacture for retaining packet order in multiprocessor systems utilizing multiple transmit queues while minimizing spinlocks are disclosed herein. Embodiments of the present invention define multiple transmit queues for a given priority level of packets to allow parallel processing and queuing of packets having equal priority in different transmit queues. Queuing packets of equal priority in different transmit queues minimizes processor time spent attempting to acquire queue-specific resources associated with one particular transmit queue. In addition, embodiments of the present invention provide an assignment mechanism to maximize utilization of the multiple transmit queues by queuing packets corresponding to each transmit request in a next available transmit queue defined for a given priority level. Coordination between hardware and software allows the order of the queued packets to be maintained in the transmission process.

(21) Appl. No.: 10/143,466

(22) Filed: May 9, 2002

Related U.S. Application Data

(63) Continuation-in-part of application No. 09/963,284,
filed on Sep. 25, 2001.

**FIG. 1**

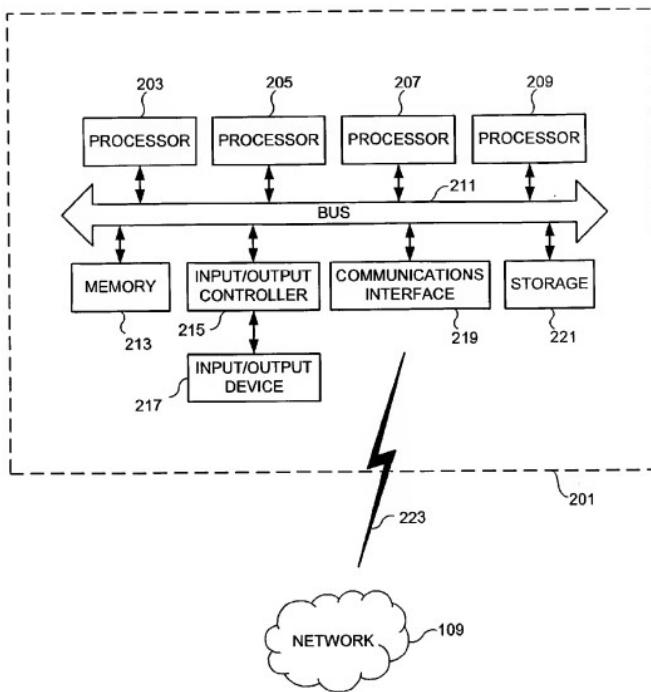
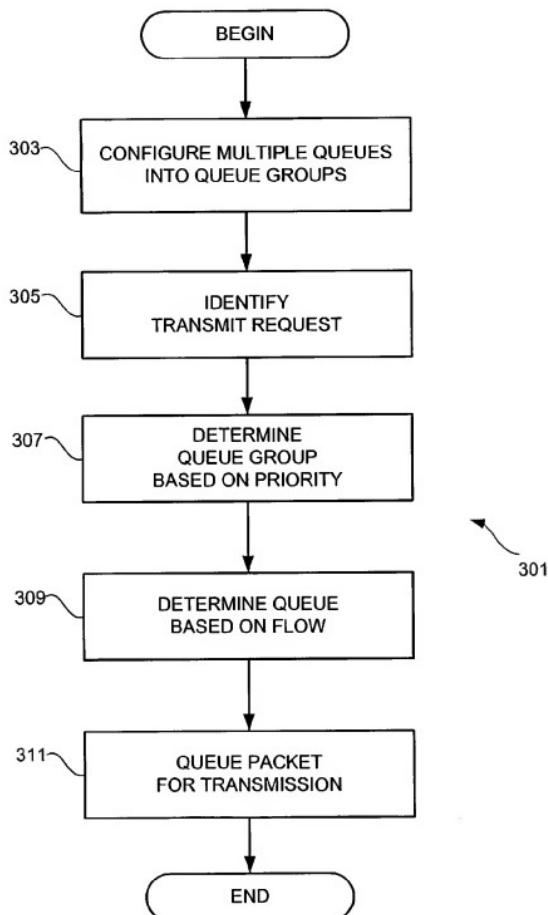


FIG. 2

**FIG. 3**

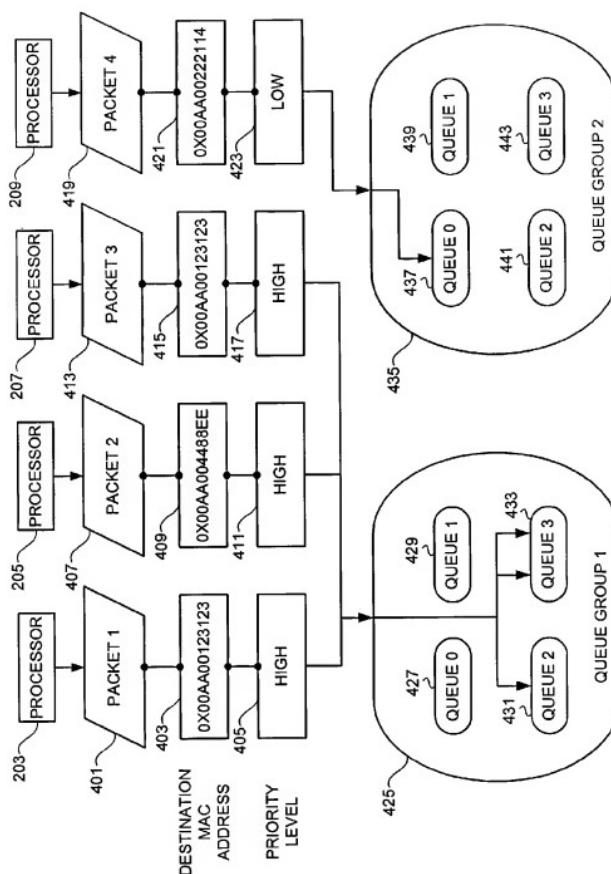
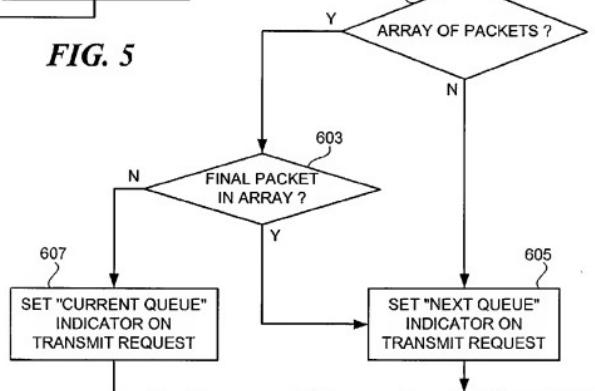
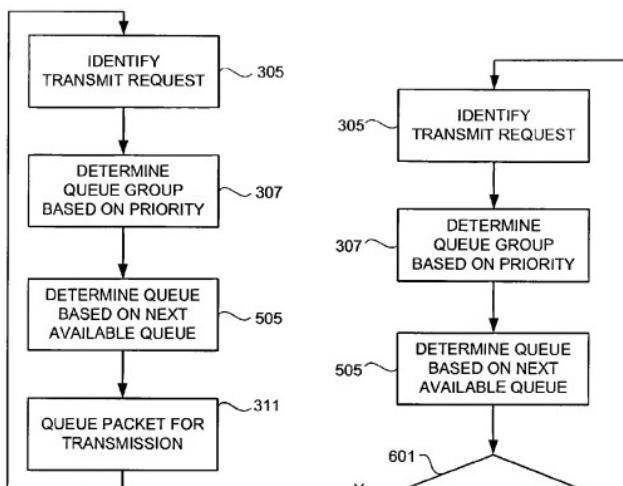


FIG. 4



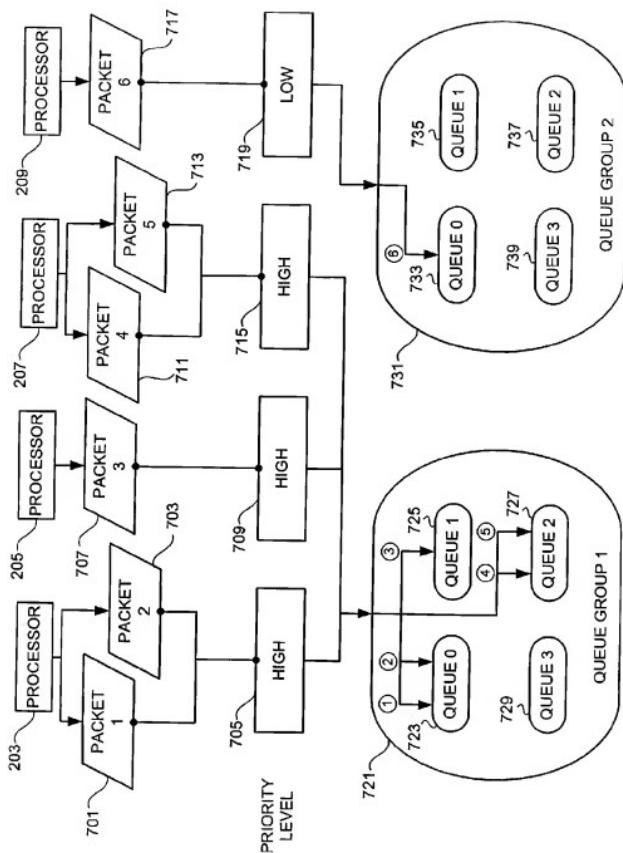


FIG. 7

METHOD AND APPARATUS FOR MINIMIZING SPINLOCKS AND RETAINING PACKET ORDER IN SYSTEMS UTILIZING MULTIPLE TRANSMIT QUEUES**CROSS-REFERENCE TO RELATED APPLICATION**

[0001] This application is a continuation-in-part of co-pending U.S. application Ser. No. 09/963,284, filed Sep. 25, 2001, and titled METHODS AND APPARATUS FOR RETAINING PACKET ORDER IN SYSTEMS UTILIZING MULTIPLE TRANSMIT QUEUES. The benefit of the filing date of the above-referenced application is hereby claimed pursuant to 35 U.S.C. §120.

TECHNICAL FIELD

[0002] This disclosure relates to network communications, and more particularly, to but not exclusively, to a method, apparatus, and article of manufacture for processing multiple transmit requests in parallel while retaining packet order within a flow to reduce the potential that packets may be received out-of-order by a destination node on a network, and consequently lost.

BACKGROUND INFORMATION

[0003] Many modern network adapters are configured to include multiple transmit queues. Typically, multiple transmit queues are made available in order to provide enhanced quality of service ("QoS"), the preferential treatment of some packets over others, by providing queues of varying priority levels (e.g., one queue per priority level), into which packets may be deposited while awaiting transmission on a network. Many network communication protocols, such as Internet Protocol, version 6 ("IPv6") (IETF draft standard RFC 2460, December 1998, "Internet Protocol, Version 6 (IPv6) Specification") provide a capability to specify a priority for the packet as a part of an IP header. Providing queues of varying priority levels enables the network adapter to transmit higher priority packets before lower priority packets, even though the higher priority packets may have been queued subsequent to the lower priority packets.

[0004] Because packets within the same flow (a single communication may comprise one or multiple packets depending on the size of the communication) will be designated with the same priority, providing one queue per priority level ensures that packets will be sent in the same order in which they were processed by a protocol stack, thereby eliminating, to the extent possible, the potential that the packets will be received out-of-order at the destination node on the network. While not an issue in regard to all communication protocols, the order in which packets are transmitted and received can represent a problem for some communication protocols and particularly one widely used communication protocol, Transmission Control Protocol ("TCP") (IETF standard RFC 793, Sep. 1, 1981, "Transmission Control Protocol"). In one implementation of TCP for example, if a packet is received out-of-order (i.e., the order in which it was processed by the protocol stack) by more than a small margin (e.g., a swap with an adjacent packet), it will be considered lost, and will have to be retransmitted. Providing one queue per one or more priority

levels helps to alleviate this problem because all packets of equal priority will be deposited in the same queue. However, in network environments in which a protocol that does not include a priority indicating feature is in use, or the protocol in use supports fewer levels of priority than there are queues available with a network adapter, at least a portion of the multiple transmit queues associated with the network adapter remain unused.

[0005] The reader will appreciate that a packet comprises a package of information transmitted as a single unit from a source node on a network to a destination node on the network, and typically includes data, as well as addressing, routability, and reliability information in the form of headers and/or footers to enable the network's communication protocols to deliver the data within the packet to the intended recipient, and to ensure that the data arrives intact and without corruption. As transmit requests are generated from applications in response to user input or otherwise, the transmit requests are transferred to a protocol stack, such as a Transmission Control Protocol/Internet Protocol ("TCP/IP") protocol stack for encapsulation into the packet. After incorporating the data and other information into the packet, the packet is transferred to a device driver associated with a network adapter to be placed in a queue, awaiting transmission on the network. The device driver generally comprises a software component that permits a computer system to communicate with a device (e.g., a network adapter), and manipulates data in order to transmit data to the device (e.g., depositing packets in a queue and informing the network adapter that packets are waiting to be transmitted).

[0006] In a multiprocessor system, it is possible to have multiple transmit requests arrive at a "send packet" function of the device driver simultaneously on different processors. As mentioned previously, if the packets of the transmit requests are of equal priority, different priorities that are sharing a queue, or a priority-indicating protocol is not being used, then the packets will be deposited in the same queue via current queuing mechanisms to await transmission on the network. In order to provide access to the "send packet" function resources associated with a particular queue, the device driver may use semaphores to indicate to potential users that the "send packet" function resources associated with the particular queue are in use, thereby preventing access by more than one thread.

[0007] In a Windows® operating system environment for example, the semaphores are referred to as "spinlocks." Spinlocks guarantee that only one process at a time will be given access to a particular resource, thereby ensuring data integrity within the resource. As a consequence of the foregoing, when a second processor, for example, attempts to acquire a resource (e.g., a queue specific element of the "send packet" function) that is already being used by a first processor, for example, the second processor "spins," and waits for the spinlock to be released. While waiting for the spinlock to be released, the second processor remains in a "busy wait," and is not free to perform other tasks. Often, in network environments designed for high throughput levels (e.g., gigabit or 10 gigabit connections), the throughput level of the network connection itself may be limited by processor bandwidth, particularly, cycles spent preparing and processing network transmit requests.

**BRIEF DESCRIPTION OF THE VARIOUS
VIEWS OF THE DRAWINGS**

[0008] In the drawings, like reference numerals refer to like parts throughout the various views of the non-limiting and non-exhaustive embodiments of the present invention, and wherein:

[0009] **FIG. 1** is a block diagram illustrating one embodiment of a network environment in accordance with an embodiment of the present invention;

[0010] **FIG. 2** is a block diagram illustrating an embodiment of a computer system in accordance with an embodiment of the present invention;

[0011] **FIG. 3** is a flow diagram illustrating an example of events in one embodiment of a process in accordance with an embodiment of the present invention;

[0012] **FIG. 4** is a schematic block diagram illustrating how multiple packets may be processed and queued in one embodiment of a multiprocessor system in accordance with an embodiment of the present invention;

[0013] **FIG. 5** is a flow diagram illustrating an example flow of events in another embodiment of a process in accordance with an embodiment of the present invention;

[0014] **FIG. 6** is a flow diagram illustrating an example flow of events in yet another embodiment of a process in accordance with an embodiment of the present invention; and

[0015] **FIG. 7** is a schematic block diagram illustrating how multiple packets may be processed and queued in one embodiment of a multiprocessor system in accordance with another embodiment of the present invention.

**DETAILED DESCRIPTION OF THE
ILLUSTRATED EMBODIMENTS**

[0016] Embodiments of a method, apparatus, and article of manufacture for minimizing spinlocks and retaining packet order in systems utilizing multiple transmit queues are described in detail herein. In the following description, numerous specific details are provided, such as the identification of various system components, to provide a thorough understanding of embodiments of the invention. One skilled in the art will recognize, however, that embodiments of the invention can be practiced without one or more of the specific details, or with other methods, components, materials, etc. In still other instances, well-known structures, materials, or operations are not shown or described in detail to avoid obscuring aspects of various embodiments of the invention.

[0017] Reference throughout this specification to "one embodiment" or "an embodiment" means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present invention. Thus, the appearance of the phrases "in one embodiment" or "in an embodiment" in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[0018] As an overview, embodiments of the invention provide a method, apparatus, and article of manufacture to

define multiple queues per priority level, or priority level group, to allow parallel processing of multiple transmit requests, corresponding to packets of equal priority or to packets having priorities corresponding to a common priority level group, that may be received substantially simultaneously at different processors within a multiprocessor system. Another aspect of embodiments of the present invention includes providing an assignment mechanism, such as a hashing algorithm, to ensure that packets corresponding to the same flow are processed in series and queued in the same transmit queue to maintain packet order and reduce the potential for out-of-order packets, which may result in lost packets and a reduction in the throughput capabilities of the network link.

[0019] In still other embodiments, the assignment mechanism may comprise utilization of a next available transmit queue within a queue group in which all packets corresponding to a single transmit request may be queued for transmission. This assignment mechanism ensures that packets from a common flow will be transmitted in order, to the extent possible, by queuing all packets corresponding to a single request in a common transmit queue, while minimizing spinlocks or other processing collisions by providing the next available transmit queue for each successive transmit request. In one embodiment, hardware (e.g., a network adapter) may then be notified (e.g., via an indicator that may be set in the transmit request) regarding a sequence of transmit queues to access in order to transmit packets in the order in which they were processed. Other features of the illustrated embodiments will be apparent to the reader from the foregoing and the appended claims, and as the detailed description and discussion is read in conjunction with the accompanying drawings.

[0020] Referring now to the drawings, and in particular to **FIG. 1**, there is illustrated a network environment in accordance with an embodiment of the present invention. In one embodiment, the network environment comprises a plurality of clients 101, 103, 105, and 107 interconnected via a network 109 by a plurality of communication links 111, 113, 115, and 117, respectively. The network 109 may be any type of communications network through which a plurality of different devices may communicate, such as for example, but not limited to, a public switched telephone network ("PSTN"), an Internet, an intranet, an extranet, a wide area network ("WAN"), a local area network ("LAN"), or other network or combination of networks to enable communication among various networks, or between any of the illustrated components connected to the network(s), or other components. The communication links 111, 113, 115, and 117 between the clients 101, 103, 105, and 107, respectively, and the network 109 may comprise wires, cables, optical fibers, or other physical connections in various embodiments of the invention. In other embodiments, the communication links 111, 113, 115, and 117 may comprise a plurality of wireless links utilizing some portion of the electromagnetic spectrum, such as for example, but not limited to, radio frequency or infrared signals. In still other embodiments, the communication links 111, 113, 115, and 117 may comprise an optical link, or any combination of the foregoing.

[0021] **FIG. 2** is a block diagram illustrating one embodiment of a machine 201 that may be used for the clients 101, 103, 105, and 107 in accordance with an embodiment of the present invention. Typically, the clients 101, 103, 105, and

107 may comprise various types of machines, including a desktop computer or a workstation, a laptop computer, a personal computer, or the like. In one embodiment, the machine 201 is a computer system that includes a plurality of processors 203, 205, 207, and 209 coupled to a bus 211. In one embodiment, a memory 213, a storage 221, a communications interface 219, and an input/output controller 215 are also coupled to the bus 207.

[0022] In one embodiment, the machine 201 interfaces to external systems, such as the network 109, through the communications interface 219. The communications interface 219 may include a radio transceiver compatible with various modulated signals, wireless telephone signals, or the like. The communications interface 219 may also include an Ethernet adapter, an analog modem, Integrated Services Digital Network ("ISDN") modem, cable modem, Digital Subscriber Line ("DSL") modem, a T-1 line interface, a T-3 line interface, an optical carrier interface (e.g., OC-3), token ring interface, satellite transmission interface, a wireless interface, or other interfaces for coupling a device to other devices.

[0023] In one embodiment, a carrier wave signal 223 is received/transmitted between the communications interface 219 and the network 109. In one embodiment, the communications signal 223 may be used to interface the machine 201 with another computer system, a network hub, a router, or the like. In one embodiment, the carrier wave signal 223 is considered to be machine-readable media, which may be transmitted through wires, cables, optical fibers, or through the atmosphere, or the like.

[0024] In one embodiment, the processors 203, 205, 207, and 209 may be conventional processors, such as for example, but not limited to, an Intel x86 processor, or Pentium family microprocessor, a Motorola family microprocessor, or the like. The memory 213 may be a machine-readable medium (also "machine-accessible medium") such as dynamic random access memory ("DRAM"), and may include static random access memory ("SRAM"). An input/output device 217, coupled to the input/output controller 215 may be a keyboard, a disk drive, a printer, a scanner, or other input/output device, including a television remote, a mouse, a trackball, a trackpad, a joystick, or the like.

[0025] The storage 221, in one embodiment, may include machine-readable media such as, for example, but not limited to, a magnetic hard disk, a floppy disk, an optical disk, a read-only memory component ("ROM"), a smart card, or another form of storage for data. In one embodiment, the storage 221 may include removable media, read-only memory, readable/writable memory, or the like. Some of the data may be written by a direct memory access process into the memory 213 during execution of software in the computer system 201. It will be appreciated that software may reside in the storage 221, the memory 213, or may be transmitted or received via a modem or the communications interface 219. For the purpose of the specification, the terms "machine-readable medium" and "machine-accessible medium" shall be taken to include any medium that is capable of storing data, information, or encoding a sequence of instructions or operations for execution by the processors 203, 205, 207, and 209 to cause the processors 203, 205, 207, and 209 to perform the methodologies in accordance with an embodiment of the present invention. The terms

"machine-readable medium" and "machine-accessible medium" shall be understood to include, for example, solid-state memories; ROM; random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier tones, infrared signals, and digital signals); and the like.

[0026] It will be appreciated that other components may be included or substituted for those shown in the illustrated embodiment. For example, the machine 201 illustrated in FIG. 2 may include a greater or lesser number of processors (such as, e.g., processors 203, 205, 207, and 209), and may include other components such as a display controller, an audio controller, or the like.

[0027] With continued reference now primarily to FIGS. 1 and 2, in one embodiment of the present invention, the communications interface 219 comprises a network adapter configured to provide multiple transmit queues for queuing packets awaiting transmission on the network 109. It will be appreciated that in one embodiment, the multiple transmit queues comprise memory allocations within the memory 213, for example, and are not physically a part of the communications interface 219. In one embodiment, the communications interface 219 includes a communications memory (not shown), which may include a plurality of registers to inform the communications interface 219 of the location of the memory allocations comprising the multiple transmit queues.

[0028] In accordance with an embodiment of the present invention, the multiple transmit queues may be configured, by a device driver at system initialization for example, to define at least two queues for a given priority level, or group of priority levels ("priority level group") supported by a network communication protocol. The priority level group may comprise a single priority level in one embodiment of the present invention. In other embodiments, the priority level group may comprise two or more priority levels. The network communication protocol may correspond to a protocol (e.g., IPv6) being utilized by the plurality of clients 101, 103, 105, and 107 to communicate information on the network 109, in an embodiment. It will be appreciated that in circumstances in which the communication protocol being used by the clients 101, 103, 105, and 107 to communicate information on the network 109 does not support the designation of priority levels for packets, all available transmit queues may be defined for the same "priority" (e.g., none will have a higher or lower transmission priority than any other), in an embodiment. For example, if a network adapter provides for sixteen transmit queues and the communications protocol provides eight levels of priority, then two transmit queues may be defined for each level of priority, in one embodiment. Defining multiple transmit queues per priority level, or priority level group, permits the multi-processor system (e.g., the machine 201 illustrated in FIG. 2) to process transmit requests, including those corresponding to packets of equal priority, or those having priorities corresponding to a common priority level group, substantially in parallel, and thereby alleviate any network throughput limitations or wasted CPU cycles associated with the processing of transmit requests of the same priority simultaneously. It will be appreciated that the ability to process transmit requests in parallel is particularly relevant in situations in which a communications protocol having no

priority-indicating feature is in use because all packets would, without reference to embodiments the present invention, be queued in the same transmit queue, thereby preventing utilization of the multiple processors to increase a rate at which transmit requests are processed.

[0029] With reference now primarily to FIG. 3, an example flow of events in the implementation of a process 301 is shown in accordance with an embodiment of the present invention. In one instance, the flow of events illustrated in FIG. 3 may be embodied in the form of a set of machine-readable or machine-accessible instructions comprising a part of the device driver discussed above. In one embodiment, the device driver may be stored in the memory 213, the storage 221, or received by the machine 201 via the communications interface 219 (see, e.g., FIG. 2). The process 301 begins with a configuration of the multiple transmit queues provided by the communications interface 219 by defining a plurality (e.g., two or more) of transmit queues per priority level, or group of priority levels, (see, e.g., process block 303) to generate one or more queue groups (e.g., a group of transmit queues defined to correspond to a particular priority level or set of priority levels). It will be appreciated that in various embodiments of the present invention, multiple transmit queues may be defined for one or more priority levels, while defining only a single queue for other priority levels.

[0030] As transmit requests, each corresponding to a packet, for example, are received by the "send packet" function of the device driver on a plurality of processors (e.g., processors 203, 205, 207, and 209) in a multiprocessor system (see, e.g., process block 305), the queue group corresponding to the packet's priority level may be determined (see, e.g., process block 307). Using the example above wherein the network adapter provides for sixteen transmit queues, and the communications protocol provides eight levels of priority, a total of eight queue groups, each comprising two transmit queues, may be provided, each queue group corresponding to a particular priority level designation, or set of priority level designations. Following a determination of the appropriate queue group based on the packet's priority designation, the process 301 proceeds to determine the appropriate transmit queue within the queue group in which to queue the packet to await transmission on the network (see, e.g., process block 309).

[0031] As mentioned previously, some popular transmission protocols (e.g., TCP) are designed such that the packet order may be important, and should therefore be considered in a determination of which transmit queue is appropriate for a particular packet. The theory being, that if all packets from a given flow are queued in the same transmit queue, the packet order will be preserved and the likelihood of dropped packets will be minimized. Since packets corresponding to a common flow will be designated with the same priority level, they are already destined to be queued within the same queue group. A determination of the specific transmit queue within a given queue group may be based on an identifying characteristic associated with the packet, preferably a characteristic common to a given flow of packets, in an embodiment.

[0032] For example, in accordance with one embodiment of the present invention, a destination media access control ("MAC") address may be used as the identifying character-

istic for selecting a transmit queue within a given queue group. With respect to the example given above wherein each queue group includes two transmit queues, a hashing algorithm for example, may consider the last bit of the destination MAC address to determine which transmit queue of the two available transmit queues within the queue group is appropriate. Table I, below, lists the possible hexadecimal digits that may comprise the destination MAC address in one embodiment, and the value corresponding to each of the four bits representing each hexadecimal digit. As the reader will appreciate, the value of the last bit will always be either "0" or "1," and may be defined to correspond to the two available transmit queues within the queue group, respectively. Since the destination MAC will be identical for all packets corresponding to a common flow, all packets from the flow will be queued in the same transmit queue within the queue group, thereby maintaining packet order. It will be appreciated that other identifying characteristics may also be determined the specific transmit queue for a given packet.

TABLE I

Hexadecimal	Binary
0	0000
1	0001
2	0010
3	0011
4	0100
5	0101
6	0110
7	0111
8	1000
9	1001
A	1010
B	1011
C	1100
D	1101
E	1110
F	1111
Value:	8421

[0033] After the specific transmit queue, within the queue group, has been determined (see, e.g., black 309), then the packet may be queued to await transmission on the network (see, e.g., process block 311) by accessing the queue-specific resources of the "send packet" function of the device driver.

[0034] It will be appreciated that a greater number of transmit queues may also be defined for a given priority level, or that different numbers of transmit queues may be defined for different priority levels based on anticipated transmit requests for a given priority level, or to ensure that higher priority packets are being transmitted as quickly as possible (e.g., define a greater number of transmit queues for higher priority levels than lower priority levels). In these various scenarios, the hashing algorithm discussed above may consider a greater number of bits of the destination MAC address for example. In one instance, if four transmit queues are defined for a given priority level, then consideration of the last two bits of the destination MAC address will hash to four different possible values (e.g., 0, 1, 2, and 3), which may be defined to correspond to the four available transmit queues. If all sixteen available transmit queues are defined for a single priority level (e.g., a non-priority indicating protocol is in use), then consideration of the last four bits of the destination MAC address will hash to sixteen possible values (e.g., 0-15), which may be defined to cor-

respond to the sixteen available transmit queues corresponding to that priority level or queue group. One skilled in the art will recognize that various other hashing, lookup, or other algorithms, that may encompass one or more portions of the packets, such as the destination IP address or TCP port, in addition to or instead of the destination MAC address, may be used in various embodiments of the present invention.

[0035] With reference now primarily to FIG. 4, a pictorial representation of how packets may be processed and queued in an embodiment of the multiprocessor system is shown in accordance with an embodiment of the present invention. For purposes of illustrating embodiments of the present invention, assume that the client 101 (see, e.g., FIG. 1) is configured in a manner such as that illustrated in FIG. 2 with four processors 203, 205, 207, and 209 capable of receiving transmit requests generated from applications, user input, or otherwise. Assume also that the communications interface 219 (see, e.g., FIG. 2) is designed to provide a total of eight transmit queues 427, 429, 431, 433, 437, 439, 441, and 443, which, at system initialization, are configured by the device driver to define four transmit queues for a "high" priority level queue group 425 (e.g., queue group 1, or "the first queue group") and a "low" priority level queue group 435 (e.g., queue group 2, or "the second queue group"), respectively, for example. The two priority levels may correspond to priority levels associated with a communications protocol, for example, or may correspond to a system-specific configuration, or the like.

[0036] With continued reference to the embodiment illustrated in FIG. 4, a transmit request, corresponding to a packet, is received at each of the four processors 203, 205, 207, and 209 substantially simultaneously. A first packet 401, including a destination MAC address 403 and a priority designation 405, is received at the first processor 203. A second packet 407, including a destination MAC address 409 and a priority designation 411, is received at the second processor 205. A third packet 413, including a destination MAC address 415 and a priority designation 417, is received at the third processor 207, and a fourth packet 419, including a destination MAC address 421 and a priority designation 423, is received at the fourth processor 209. It will be appreciated that in another embodiment of the present invention, the client 101 (see, e.g., FIG. 1) may include only a single processor with multiple threads, thereby facilitating the receipt of multiple transmit requests substantially simultaneously.

[0037] As mentioned above in conjunction with FIG. 3, the appropriate queue group (e.g., based on the defined priority levels) may be determined with reference to the respective priority designation given to a particular packet. For example, the first packet 401 has a priority designation 405 of "High" and is consequently placed in the first queue group 425. Similarly, the second packet 407 and the third packet 413 also have priority designations 411, and 417, respectively, of "High" and are also assigned to the first queue group 425. The fourth packet 419, on the other hand, has a priority designation 423 of "Low" and is consequently assigned to the second queue group 435. Determining the specific transmit queue within the queue group then becomes a matter of hashing the destination MAC address to a value corresponding to one of the available transmit queues 427, 429, 431, 433, 437, 439, 441, and 443 within the

respective queue groups 425 and 435, in an embodiment. Using the example described above in conjunction with FIG. 3, where four transmit queues are available in a given queue group, such as they are in the example illustrated in FIG. 4, the last two bits of the destination MAC address may be used to assign the packet to a specific transmit queue. The first packet 401, with a destination MAC address 403 ending with the hexadecimal digit "3," includes "0011" as the final four bits of the address 403. Using the values assigned to each of the last two bits (e.g., "11") in Table 1, a value of 3 may be calculated (e.g., one for the last bit, plus two for the preceding bit), which may be defined to correspond to a third transmit queue 433 within the first queue group 425 in an embodiment.

[0038] Having determined which transmit queue (e.g., the third transmit queue 433) the first packet 401 should be queued in to await transmission on the network 109 (see, e.g., FIG. 1), the processor 203 must acquire the queue-specific resources of the device driver's "send packet" function associated with the third transmit queue 433. In this example, the queue-specific resources of the "send packet" function associated with the third transmit queue 433 are not currently being used and consequently, the first packet 401 may be queued in the third transmit queue 433 so that it may be accessed by the communications interface 219 (see, e.g., FIG. 2) and transmitted on the network 109 in due course.

[0039] Concurrently, while the first packet 401 is being processed by the first processor 203, the second packet 407 may be processed by the second processor 205. The second packet 407 includes a destination MAC address 409 ending with the hexadecimal digit "E," which corresponds to the four bits "1110" (see, e.g., Table 1). Using the last two bits (e.g., "10"), and the values associated with them in Table 1, a value of 2 may be calculated, which, in the illustrated embodiment, corresponds to a second transmit queue 431 within the first queue group 425. As with the first processor 203, the second processor, having determined the appropriate transmit queue (e.g., the second transmit queue 431) for the second packet 407, will attempt to acquire the queue-specific resources of the device driver's "send packet" function associated with the second transmit queue 431. Since no other processor is currently utilizing these resources, the second processor 205 may acquire them and queue the second packet 407 in the second transmit queue 431.

[0040] Processing and queuing these two packets of equal priority in parallel allows the multiprocessor machine 201 (see, e.g., FIG. 2) to function more efficiently by decreasing acquisition time of queue-specific resources. By permitting packets of equal priority to be queued in different transmit queues, instances of processor "busy waits," while attempting to acquire a spinlock, may be reduced. It will be appreciated however, that equal priority packets from different flows may also be queued in the same transmit queue, depending on the calculated value from the hashing algorithm, or other queue-assignment mechanism.

[0041] Continuing with the foregoing example, the third packet 413, with a destination MAC address 415 ending with the hexadecimal digit "3," also corresponds to the third transmit queue 433 within the first queue group 425. The reader will note that the destination MAC address 415 of the third packet 413 is identical to the destination MAC address

[0043] of the first packet 401, indicating that both are destined to the same client (e.g., the client 103 of FIG. 1) on the network 109, and should therefore be queued in the same transmit queue in order to prevent transmission of out-of-order packets. In this case, the third processor 207, having determined that the third packet 413 corresponds to the third transmit queue 433, will attempt to acquire the queue-specific resources of the device driver's "send packet" function associated with the third transmit queue 433. Because these resources are already in use by the first processor 203, the third processor 207 will not be permitted to acquire the resources, and will have to "spin" and wait for the resources to be released by the first processor 203. This ensures that the packets will be deposited in the transmit queue and transmitted in order to enable them to be received in order, to the extent possible, by the destination client on the network 109.

[0042] In a manner similar to that described above, the fourth packet 419 will be processed to determine the appropriate transmit queue within the second queue group 435 by evaluating the last two bits of the destination MAC address 421, namely "00" (corresponding to "0100," the four bits corresponding to the hexadecimal digit "4"). In this case, the last two bits hash to a value of "0," which may be defined to correspond to the zero transmit queue 437 within the second queue group 435. The fourth processor 209 will then acquire the resources associated with the zero transmit queue 437, as above, and queue the fourth packet 419 to await transmission on the network 109. The reader will appreciate that any number of transmission scenarios may be used to actually send the queued packets. In one scenario, the communications interface 219 may be configured to access the four transmit queues 427, 429, 431, and 433 of the first queue group 425 in sequence until all high priority packets have been transmitted, and only then access the transmit queues 437, 439, 441, and 443 of the second queue group 435 in order to transmit the lower priority packets on the network 109. In another scenario, two packets may be transmitted from high-priority transmit queues for every one packet transmitted from a low-priority transmit queue, for example.

[0043] In another embodiment of the present invention, the assignment mechanism for determining the appropriate transmit queue within a given queue group, in which a packet should be queued, may be based on a "next available" standard. For example, each transmit request may cause the processor at which it is received to acquire the next available queue-specific resources in order to process the packet(s) associated with the transmit request. This assignment mechanism may be particularly appropriate under conditions in which a large number of packets are being communicated via a proxy server, or to destinations on an opposite side of a router, or the like. In these circumstances, assigning packets to specified queues in response to a hashing algorithm or the like that considers an attribute of the packet (e.g., one or more bits of the destination MAC address) in making the assignment, may result in the large number of packets being queued in a common queue, thereby reducing the efficiency by which packets may be processed in a multiprocessor system. By assigning packets to a queue based on availability, the system may make more efficient use of the multiple queues defined for a priority level queue group and minimize spinlock collisions as multiple transmit requests are processed in parallel.

[0044] Depending on the operating system and the communication protocol being used, transmit requests may comprise only a single packet, or may comprise one or multiple packets in response to the size of the communication. While a single process may be implemented to function under either of these conditions (see, e.g., FIG. 6), a less complex process may be utilized in those situations in which each transmit request is known to correspond to a single packet (see, e.g., FIG. 5). By coordinating the order in which packets may be queued in the transmit queues with the order in which they may be processed by the hardware for transmission, absolute ordering may be preserved.

[0045] With reference now primarily to FIG. 5, a flow diagram illustrating an example flow of events in a process for queuing packets is shown in accordance with an embodiment of the present invention. The process illustrated in FIG. 5 may be implemented under conditions in which each transmit request is known to correspond to a single packet. In a manner similar to that described above in conjunction with FIG. 3, the process illustrated in FIG. 5 begins with the identification of transmit requests (see, e.g., process block 305) as they are received by the "send packet" function of the device driver on a plurality of processors (e.g., processors 203, 205, 207, and 209) in a multiprocessor system (or via multiple threads associated with a single processor). As the transmit requests are received, the queue group corresponding to the priority level of the packet associated with each transmit request may be determined (see, e.g., process block 307), as described previously.

[0046] Having determined the appropriate queue group based on the packet's priority, a specific queue within the queue group may be determined based on the next available queue (see, e.g., process block 503), in an embodiment. For instance, using the example given above wherein the network adapter provides sixteen transmit queues, and the communications protocol provides eight levels of priority, a total of eight queue groups, each comprising two transmit queues, may be provided, each queue group corresponding to a particular priority level designation, or set of priority level designations. If two transmit requests, each corresponding to a packet having the same priority designation, are received substantially simultaneously on two separate processors, a packet corresponding to a first one of the two transmit requests may be assigned to a first one of the two transmit queues. A first processor may then acquire the queue-specific resources associated with the first one of the two transmit queues corresponding to the queue group defined to correspond to the priority designation that corresponds to the example packets, and queue the packet corresponding to the first one of the two transmit requests.

[0047] In one embodiment, a packet corresponding to a second one of the two transmit requests may then be assigned to a second one of the at least two transmit queues (i.e., the next available queue). A second processor may then proceed to acquire the queue-specific resources of the second one of the two transmit queues corresponding to the queue group. Thus, instead of waiting for the first processor to finish queuing the packet corresponding to its transmit request and release the queue-specific resources associated with the first one of the two transmit queues, the second processor may process the second one of the two transmit requests in parallel with the first processor's action on the first one of the two transmit requests. By queuing each

packet in the next available queue (see, e.g., process block 311), busy waits may be reduced and the full range of available transmit queues within the queue group may be utilized, thereby enhancing the efficiency with which transmit requests are processed. As will be appreciated, the process illustrated in FIG. 5 may then repeat itself beginning again with process block 305 as additional transmit requests are received for processing.

[0048] When the time comes to transmit the packets via a communication network (e.g., the network 109, FIG. 1), the network adapter (e.g., the communications interface 219, FIG. 2) may access the queues within the queue group in the same sequence in which the queues were assigned, thereby maintaining packet order as the packets are transmitted. For example, the network adapter may begin by transmitting a packet from the first one of the two transmit queues described above in conjunction with the example wherein each queue group comprises two transmit queues. The network adapter may then transmit a packet from the second one of the two transmit queues, and then return again to the first one of the two transmit queues, and so on. Because the "send packet" function of the device driver can assign queues in an order known to (or accessible to) the hardware, absolute ordering may be preserved upon transmission.

[0049] With reference now primarily to FIG. 6, a flow diagram illustrating an example flow of events in a process for queuing packets is shown in accordance with yet another embodiment of the present invention. As mentioned previously, under conditions in which multiple packets may be associated with a single transmit request, a more detailed structure for maintaining packet order may be required in conjunction with an assignment mechanism based on the "next available" queue.

[0050] In one embodiment, the process illustrated in FIG. 6 may begin, as described previously, with the identification of transmit requests (see, e.g., process block 305) as they are received by the "send packet" function of the device driver on a plurality of processors in a multiprocessor system, or via a plurality of threads associated with a single processor. Again, as described previously, a queue group may then be determined based on the priority designation of the packet(s) associated with each transmit request (see, e.g., process block 307). It will be appreciated that if a transmit request corresponds to two or more packets, all packets corresponding to the transmit request will have the same priority, in an embodiment.

[0051] As described above in conjunction with FIG. 5, the specific queue within the queue group into which the packet(s) may be deposited to await transmission may then be determined based on the next available queue (see, e.g., process block 305). Under conditions in which multiple packets may be associated with a single transmit request, absolute ordering of packets as received from the protocol stack may be preserved by effectively "reserving" the "next available queue" for each packet in the array. This may be accomplished, in one embodiment, by queuing all packets corresponding to a single transmit request in a common transmit queue. Ensuring that this ordering is preserved during the transmission process may then be facilitated by indicating, to hardware for example, a queue location corresponding to a next packet to be transmitted. In the case of packet arrays, the queue location may comprise the same

queue as the packet currently being transmitted (or previously transmitted), while in the case of single packets per transmit request (or the final packet in an array), the queue location may comprise the next transmit queue in the sequence in which queues were assigned, as described previously in conjunction with FIG. 5.

[0052] In one embodiment, indicating the queue location corresponding to the next packet to be transmitted may include setting an indicator associated with each packet to identify the transmit queue corresponding to the next packet. For example, the process illustrated in FIG. 6 may determine whether the packet being queued corresponds to an array of packets (see, e.g., process block 601). If the packet does not correspond to an array of packets (i.e., multiple packets per transmit request), then a "next queue" indicator may be set on the transmit request (see, e.g., process block 605) to indicate to the hardware (e.g., a network adapter) that it may proceed to the next queue in a sequence to find the next packet for transmission, in a manner similar to that described above in conjunction with FIG. 5.

[0053] If the packet being queued does correspond to an array of packets (see, e.g., block 601), then the process illustrated in FIG. 6 proceeds to determine whether the packet being queued represents the final packet in the array (see, e.g., process block 603). If the packet does represent the final packet in the array then the process proceeds to set the "next queue" indicator on the transmit request (see, e.g., block 605), as described above, as if the transmit request corresponds to only a single packet. If however, the packet being queued does not represent the final packet in the array (see, e.g., block 603), then the process proceeds to set a "current queue" indicator on the transmit request (see, e.g., process block 607) to indicate to the hardware (e.g., the network adapter) that it should find the next packet for transmission in the same queue. In one embodiment, setting the "next queue" indicator and/or the "current queue" indicator may be implemented using a single bit.

[0054] It will be appreciated that setting the "next queue" indicator or the "current queue" indicator is merely one example of several possible mechanisms for accomplishing the task of informing the hardware of the appropriate location in which to find the "next" packet for transmission on the network (e.g., the network 109, FIG. 1). It will further be appreciated that reference to setting an indicator may refer to inaction. For example, a single indicator may be set to indicate that a "current queue" should be accessed again to select the next packet, while the same indicator may be left unset to indicate that a "next queue" should be accessed to locate the next packet for transmission. In another embodiment, the indicator may specify the queue (e.g., by a number corresponding to each queue) in which to find the next packet for transmission.

[0055] With reference now primarily to FIG. 7, a schematic representation of how packets may be processed and queued in an embodiment of a multiprocessor system is shown in accordance with another embodiment of the present invention. For purposes of illustrating embodiments of the present invention, assume that, like described previously in conjunction with FIG. 4, the client 101 (see, e.g., FIG. 1) is configured in a manner such as that illustrated in FIG. 2 with four processors 203, 205, 207, and 209 capable to receive transmit requests generated from applications,

user input, or otherwise. Assume also that the communications interface 219 (see, e.g., FIG. 2) is designed to provide a total of eight transmit queues 723, 725, 727, 729, 733, 735, 737, and 739, which, at system initialization, are configured by the device driver to define four transmit queues for a "high" priority level queue group 721 (e.g., queue group 1, or "the first queue group") and a "low" priority queue group 731 (e.g., queue group 2, or "the second queue group"), respectively, for example. The two priority levels may correspond to priority levels associated with a communication protocol, for example, or may correspond to a system-specific configuration, or the like.

[0056] With continued reference to the embodiment illustrated in FIG. 7, a transmit request, corresponding to at least one packet, may be received at each of the four processors 203, 205, 207, and 209 substantially simultaneously. In the illustrated embodiment, a first transmit request, corresponding to a first packet 701 and a second packet 703 each having a priority designation 705, may be received at the first processor 203. A second transmit request, corresponding to a third packet 707 having a priority designation 709, may be received at the second processor 205. Similarly, a third transmit request, corresponding to a fourth packet 711 and a fifth packet 713 each having a priority designation 715, may be received at the third processor, while a fourth transmit request, corresponding to a sixth packet 717 having a priority designation 719, may be received at the fourth processor 209. As with FIG. 4, it will be appreciated that in another embodiment of the present invention, the client 101 (sec, e.g., FIG. 1) may include only a single processor with multiple threads, thereby facilitating the receipt of multiple transmit requests substantially simultaneously.

[0057] As mentioned above in conjunction with FIGS. 5 and 6, the appropriate queue group (e.g., based on the defined priority levels) may be determined with reference to the respective priority designation given to a particular packet. For example, packets 1-5 (having reference numerals 701, 703, 707, 711, and 713) all have a priority designation of "High" and may consequently be destined for the first queue group 721, while the sixth packet 717 has a priority designation of "Low" and may consequently be destined for the second queue group 731. In one embodiment, determining the specific transmit queue within the queue group then becomes a matter of availability.

[0058] In one embodiment, the "send packet" function of the device driver may, by default, assign the at least one packet associated with each transmit request to the transmit queue of the queue group in a defined order. For example, in the embodiment illustrated in FIG. 7, the first and second packets 701 and 703, corresponding to the first transmit request, may be assigned to the zero queue 723 of the first queue group 721, and the third packet 707, corresponding to the second transmit request, may be assigned to the first queue 725 of the first queue group 721. Similarly, the fourth and fifth packets 711 and 713, corresponding to the third transmit request, may be assigned to the second queue 727 of the first queue group 721, while the sixth packet 717, corresponding to the fourth transmit request, may be assigned to the zero queue 733 of the second queue group 731 in response to the priority designation 719 associated with the sixth packet 717, and so on.

[0059] In one embodiment, each of the processors 203, 205, 207, and 209 may then acquire the queue-specific

resources associated with the assigned queue, and proceed to process and queue the packet(s) corresponding to the respective received transmit request. For example, the first processor 203, upon receiving the first transmit request, may acquire the queue-specific resources of the "send packet" function associated with the zero queue 723 in the first queue group 721, and process and queue the first packet 701 in the zero queue 723. In addition, in one embodiment, since the first packet 701 corresponds to an array of packets (i.e., multiple packets corresponding to a single transmit request), but does not represent the final packet in the array (e.g., packet 2 of 2), the processor may set the "current queue" indicator associated with the first packet 701 to indicate to the hardware (e.g., the network adapter) that it should look for the next packet to send in the same queue (e.g., the zero queue). The first processor may then also process and queue the second packet 703 in the zero queue 723 of the first queue group 721 in a similar manner, and set the "next queue" indicator associated with the second packet 703 to indicate to the hardware that it should look in the next queue (e.g., the first queue 725, corresponding to the defined sequence) for the next packet to send.

[0060] Concurrently, while the first and second packets 701, and 703 are being processed and queued by the first processor 203, the third packet 707 may be processed and queued by the second processor 205. As with the first processor 203, the second processor 205 may acquire the queue-specific resources of the "send packet" function associated with the first queue 725 of the first queue group 721, in an embodiment. The second processor 205 may then process and queue the third packet 707 in the first queue 725 of the first queue group 721 in parallel with the first processor, without having to wait for the first processor to release the queue-specific resources associated with the zero queue 723. As mentioned previously, this may contribute to a greater efficiency in processing, particularly in those circumstances in which an assignment mechanism based on a packet attribute (e.g., one or more bits of the destination MAC address) may otherwise be used to assign packets to specific queues within a queue group, and in which a proxy server, router, or the like, represents the "destination" of the packet.

[0061] Because the third packet does not correspond to an array of packets, the "next queue" indicator associated with the third packet 707 may be set to indicate to the hardware that the next packet may be obtained from the next queue in the sequence, in an embodiment. In a manner similar to that described above, it will be appreciated that the fourth and fifth packets 711, and 713 may also be processed in parallel with the first, second, and third packets, and be queued in the second queue 727 of the first queue group 721, in an embodiment. The fourth packet 711, comprising a part of an array of packets, but not corresponding to the final packet in the array, may be queued and the "current queue" indicator associated with the fourth packet 711 may be set to indicate to the hardware to return to the second queue 727 of the first queue group 721 to find the next packet for transmission. The fifth packet 713, like the second packet 703, does correspond to the final packet in the array of packets and consequently the "next queue" indicator associated with the fifth packet 713 may be set, as described above.

[0062] Concurrently with the first, second, and third processors 203, 205, and 207, respectively, the fourth processor

209 may process the sixth packet **717** in conjunction with the fourth transmit request. Because the sixth packet **717** has a priority designation of "Low," the sixth packet **717** may be assigned to the zero queue **733** of the second queue group **731**, in an embodiment. The fourth processor **209** may acquire the queue-specific resources of the "send packet" function associated with the zero queue **733** of the second queue group **731**, and process and queue the sixth packet **717** and set the "next queue" indicator associated with the sixth queue **717** to indicate to the hardware to look in the next queue for the next packet to transmit.

[0063] It will be appreciated that any number of transmission scenarios may be used to actually send the queued packets. In one scenario, the communications interface **219** (see, e.g., FIG. 2) may be configured to access the transmit queues **723**, **725**, **727**, and **729** of the first queue group **721** in the sequence in which they were made available by the "send packet" function of the device driver. By proceeding to the next packet in response to the indicated queue location (e.g., "current queue," "next queue"), absolute order of the packets may be maintained.

[0064] While the invention is described and illustrated here in the context of a limited number of embodiments, the invention may be embodied in many forms without departing from the spirit of the essential characteristics of the invention. The illustrated and described embodiments, including what is described in the abstract of the disclosure, are therefore to be considered in all respects as illustrative and not restrictive. The scope of the invention is indicated by the appended claims rather than by the foregoing description, and all changes that come within the meaning and range of equivalency of the claims are intended to be embraced therein.

What is claimed is:

1. A method, comprising:

defining at least two transmit queues for a priority level group;

receiving at least two transmit requests substantially simultaneously, the at least two transmit requests corresponding to at least one packet, respectively, the at least one packet having a priority designation corresponding to the priority level group; and

processing the at least two transmit requests in parallel to queue the at least one packet corresponding to a first one of the at least two transmit requests in a first available one of the at least two transmit queues, and to queue the at least one packet corresponding to a second one of the at least two transmit requests in a next available one of the at least two transmit queues.

2. The method of claim 1, further comprising, indicating a queue location corresponding to a next packet to be transmitted.

3. The method of claim 1, wherein the priority level group corresponds to at least one communication protocol priority level.

4. The method of claim 1, wherein the priority level group comprises a single priority level.

5. The method of claim 1, wherein the priority level group comprises at least two priority levels.

6. The method of claim 2, wherein indicating a queue location includes setting an indicator associated with the at

least one packet to identify one of the at least two transmit queues, the identified one of the at least two transmit queues corresponding to the next packet to be transmitted.

7. A method, comprising:

defining at least two transmit queues to correspond to at least one priority level;

receiving at least two transmit requests at a send packet function of a device driver substantially simultaneously, each of the at least two transmit requests corresponding to at least one packet, respectively;

assigning the at least one packet corresponding to each of the at least two transmit requests to a queue group in response to a priority designation, the queue group including the at least two transmit queues corresponding to the at least one priority level;

queuing the at least one packet corresponding to a first one of the at least two transmit requests in a first available one of the at least two transmit queues; and

queuing the at least one packet corresponding to a second one of the at least two transmit requests in a next available one of the at least two transmit queues.

8. The method of claim 7, wherein the priority designation corresponds to a communication protocol priority level.

9. The method of claim 7, wherein assigning the at least one packet corresponding to each of the at least two transmit requests to a queue group comprises correlating the priority designation with at least one priority level associated uniquely with the queue group.

10. The method of claim 7, further comprising, indicating a queue location corresponding to a next packet to be transmitted.

11. The method of claim 10, wherein indicating a queue location includes setting an indicator associated with the at least one packet corresponding to each of the at least two transmit requests to identify one of the at least two transmit queues, the identified one of the at least two transmit queues corresponding to the next packet to be transmitted.

12. An article of manufacture, comprising:

a machine-accessible medium including data that, if accessed by a machine, cause the machine to, define at least two transmit queues for a priority level group;

receive at least two transmit requests substantially simultaneously, the at least two transmit requests corresponding to at least one packet, respectively, the at least one packet having a priority designation corresponding to the priority level group; and

process the at least two transmit requests in parallel to queue the at least one packet corresponding to a first one of the at least two transmit requests in a first available one of the at least two transmit queues, and to queue the at least one packet corresponding to a second one of the at least two transmit requests in a next available one of the at least two transmit queues.

13. The article of manufacture of claim 12, further including data that, if accessed by the machine, cause the machine to indicate a queue location corresponding to a next packet to be transmitted.

14. The article of manufacture of claim 12, wherein the priority level group corresponds to at least one communication protocol priority level.

15. The article of manufacture of claim 12, wherein the priority level group comprises a single priority level.
16. The article of manufacture of claim 12, wherein the priority level group comprises at least two priority levels.
17. The article of manufacture of claim 13, wherein the data that, if accessed by the machine, cause the machine to indicate a queue location includes data that, if accessed by the machine, cause the machine to set an indicator associated with the at least one packet to identify one of the at least two transmit queues, the identified one of the at least two transmit queues corresponding to the next packet to be transmitted.
18. A article of manufacture, comprising:
- a machine-accessible medium including data that, if accessed by a machine, cause the machine to,
 - define at least two transmit queues to correspond to at least one priority level;
 - receive at least two transmit requests at a send packet function of a device driver substantially simultaneously, each of the at least two transmit requests corresponding to at least one packet, respectively;
 - assign the at least one packet corresponding to each of the at least two transmit requests to a queue group in response to a priority designation, the queue group including the at least two transmit queues corresponding to the at least one priority level;
 - queue the at least one packet corresponding to a first one of the at least two transmit requests in a first available one of the at least two transmit queues; and
 - queue the at least one packet corresponding to a second one of the at least two transmit requests in a next available one of the at least two transmit queues.
19. The article of manufacture of claim 18, wherein the priority designation corresponds to a communication protocol priority level.
20. The article of manufacture of claim 18, wherein the data that, if accessed by the machine, cause the machine to assign the at least one packet corresponding to each of the at least two transmit requests to a queue group comprises data that, if accessed by the machine, cause the machine to correlate the priority designation with at least one priority level associated uniquely with the queue group.
21. The article of manufacture of claim 18, further including data that, if accessed by the machine, cause the machine to indicate a queue location corresponding to a next packet to be transmitted.
22. The article of manufacture of claim 21, wherein the data that, if accessed by the machine, cause the machine to indicate a queue location includes data that, if accessed by the machine, cause the machine to set an indicator associated with the at least one packet corresponding to each of the at least two transmit requests to identify one of the at least two transmit queues, the identified one of the at least two transmit queues corresponding to the next packet to be transmitted.
23. A system, comprising:
- a plurality of processors to receive at least two transmit requests substantially simultaneously, each of the at least two transmit requests corresponding to at least one packet having a priority designation corresponding to a priority level group;
 - a communications interface, communicatively coupled to the plurality of processors, to transmit packets on a network; and
 - a memory, communicatively coupled to the plurality of processors, to store data and to provide at least two transmit queues accessible to the communications interface, the at least two transmit queues defined to correspond to the priority level group, and
- in response to receipt of the at least two transmit requests substantially simultaneously at the plurality of processors, the plurality of processors process the at least two transmit requests in parallel to queue the at least one packet corresponding to a first one of the at least two transmit requests in a first available one of the at least two transmit queues and to queue the at least one packet corresponding to a second one of the at least two transmit requests in a next available one of the at least two transmit queues.
24. The system of claim 23, wherein the plurality of processors comprise a single processor with multiple threads.
25. The system of claim 23, wherein the priority level group corresponds to at least one communication protocol priority level.
26. The system of claim 23, wherein the plurality of processors indicate a queue location, corresponding to a next packet to be transmitted via the communications interface, in conjunction with each packet corresponding to each of the at least two transmit requests.
27. The system of claim 26, wherein the plurality of processors indicate the queue location by setting an indicator associated with the at least one packet to identify one of the at least two transmit queues, the identified one of the at least two transmit queues corresponding to the next packet to be transmitted via the communications interface.

* * * * *